DOCUMENT RESUME

ED 311 893 IR 014 188

AUTHOR Anbar, Michael; Anbar, Ada
TITLE The "Understanding" of Natural Language in CAI and Analogous Mental Processes.
PUB DATE Jun 88
NOTE 8p.; Paper presented at the Symposium on the Engineering of Computer-Based Medical Systems (Minneapolis, MN, June 8-10, 1988).
PUB TYPE Viewpoints (120) -- Reports - Descriptive (141) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Authoring Aids (Programing); *Clinical Diagnosis; *Cognitive Processes; Communication (Thought Transfer); *Communication Disorders; Computer Assisted Instruction; Computer Simulation; Higher Education; Medical Schools; *Microcomputers; *Models; Tests
IDENTIFIERS Clustering (Learning); *Natural Language; Parallel Processes

ABSTRACT

This paper argues that the use of unrestricted natural language (UNL) interaction is essential in computer-assisted instruction aimed at teaching clinical decision making, and in interactive tests that emulate oral examinations. Knowledge about the scope and limitations of computerized UNL interaction gained during the development of the authoring tool CASIP (Computer Assisted Socratic Instructional Program) serves as the basis for a discussion of the factors involved in teaching a computer to understand natural language, including the conceptual gap between computer languages and human natural languages, and how a computer can be programmed to understand the meaning of sentences. An analogy is then drawn between CASIP and communication between humans and the discussion focuses on three questions: (1) whether this analogy exists because CASIP has been programmed to follow the pattern of human communication; (2) whether this analogy is coincidental given the assumption that mental processes proceed in parallel while CASIP's understanding of UNL is a linear process; and (3) whether this computer-based model represents a preferred way to the understanding of natural language by both humans and machines. A discussion of the potential of the computer's UNL interactive program model for the study of impairments in human communication concludes the report. (15 references) (CGD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

THE "UNDERSTANDING" OF NATURAL LANGUAGE IN CAI AND ANALOGOUS MENTAL PROCESSES

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
X This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

Michael Anbar and Ada Anbar

School of Medicine and Biomedical Sciences,
and Faculty Educational Studies SUNY, Buffalo, NY

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Michael Anbar
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

ED311893

The use of open ended natural language interaction is essential in computer assisted instruction (CAI) aimed at teaching clinical decision making, and in interactive tests that emulate oral examinations. In developing CASIP, a CAI authoring tool, we have learned about the scope and limitations of computerized natural language interaction. Different approaches were tested and the most effective of these was found to involve a conditionally branched search of key words from among clusters of contextually synonymous words. This mode of man-machine interaction offers a multistage model for certain aspects of human verbal interaction. Ways are suggested to test this model and use it to diagnose certain cognitive dysfunctions, especially in children.

INTRODUCTION

The use of unrestricted natural language (UNL) is essential in computer assisted instruction (CAI) aimed at teaching clinical decision making, (1) and in interactive tests that emulate oral examinations. (2) Although multiple choice CAI is much easier to author using one of the many authoring languages available, (3) its limitations cannot be overemphasized. (1,4) Computerized clinical case simulations which use natural language input have been independently developed and implemented in three medical centers. (5-10) These experiences demonstrate the feasibility of routine use of UNL in CAI, and their success encourages their wider use.

In developing and implementing the authoring tool CASIP (Computer Assisted Socratic Instructional Program) at Buffalo, (7,8) we learned a great deal about the scope and limitations of computerized UNL interaction, and wish to discuss some of our findings and conclusions. We will then compare computerized natural language understanding with verbal interaction among humans. We suggest that the program by which computers can be made to logically respond to UNL may serve as a model for the corresponding processes in the human mind, i.e., finding the most effective way to communicate with each other. Although it is an oversimplification, this assumption can lead to the understanding of certain verbal communication dysfunctions, expressed in learning disabilities. We also believe that our insight may lead to an effective treatment of some of these conditions.

TEACHING A COMPUTER TO UNDERSTAND NATURAL LANGUAGE

Basic to natural language understanding by computers are the following facts. Each character in the alphabet and every reading sign has a specific numerical value in any machine language. Each word can be expressed as an array ("string") of characters. Practically every computer language has an algorithm for comparing the values of connective members of two numerical arrays, which is used to determine if two such arrays (representing words) are equal. This enables computers to recognize any sequence of characters, i.e. any word, and is the key to the development of all computer languages.

In the development of computer languages the choice of the sequence of characters that denote a command or a function is arbitrary. But, for our own mental convenience we usually choose sequences that denote terms logically associated in our natural language with whatever we want the computer to do, e.g., "include", "if", "else" "while", "print", "return" etc. To effectively communicate with the computer we do not need more than 100 such "words".

Each computer language has its own restricted vocabulary. Identical words may invoke different functions in different languages. In most programming languages one must be careful not to use any of these "reserved" words indiscriminatively to name variables, since the computer amy take such a name of a variable as a command. All computer language have in common the use of single strings as commands or functions; all their messages can be coded in the form of a limited number of single words. They do not need the use of sentences.

IR 014188

BEST COPY AVAILABLE

In addition to the basic vocabulary used for programming there is the specific vocabulary of variables used in each particular application of a given program. For instance, if the computer asks you "What is the capital of France?", the word "PARIS" (or "Paris" or "paris") will invoke a certain response, whereas any other string of letters will invoke a different response. Verbal variables may include sentences, since the space character can be included as a member of a recognized string, e.g., the computer may be programmed to respond positively only if the answer is the "The capital of France is Paris" but negatively to "Paris is the capital of France" or "France's capital is Paris". Trying to communicate with a computer in unrestricted natural language requires, therefore, much more than the simple recognition of strings of characters. (11)

When do we need to communicate with the computer in UNL? We do not need this to make the computer carry out a command, such as resetting its clock, recalculating a spreadsheet, or providing a list of drug interactions from a data base. To achieve the we can readily learn a limited vocabulary and treat the computer as we treat our dog who understands a limited number of single word commands. In fact, given all the "shift", "control" and "alt" keys available on the keyboard we can often interact effectively with the computer through single key stroke commands. Although it would be nice if the computer would not limit us to a certain vocabulary, or require us to memorize a log set of key strokes, we can readily live with this inconvenience, since the cost of using UNL is not trivial.

We must, however, use UNL when we want the computer to emulate some intellectual function, such as responding to a wish or requesting indeterminate information, e.g., "I wish people loved me, how shall I go about it?", "What would you do if you won a million dollars in the lottery?". UNL is also necessary when we use the computer to evaluate and criticize a decision we make e.g., "I want to buy this new car, what do you think?". In each of these trivial examples, just as in arriving at a clinical diagnosis or deciding on a treatment strategy, the computer cannot come up with a meaningful response without asking the user a series of questions, each depending on the answer to its predecessors, and then base its final response on all the information given by the user. In other words, we need UNL when the computer emulates a human expert. UNL is essential for realistic computerized expert systems. In most instructional clinical simulations the computer emulates the supervising attending physician. This was our reason for developing CASIP - a CAI authoring tool with the ability to respond to UNL.

There is a substantial conceptual gap between computer languages and human natural language. Not only do we use in our verbal exchange a much larger vocabulary of words and phrases (fixed sequences of words that express some specific meaning), but we put these in sentences (groups of words and phrases of unrestricted length or order). Moreover, different words or phrases may convey the same information (constituting synonyms), while the same words may convey different information, depending on other words and their order of appearance in the sentence (Contextual information), or even on the content of the questions or statements that preceded the given sentence (conversational context). Furthermore, in our daily language there are often words which are redundant and convey no contextual information. The same words may, however, convey meaningful information in other contexts. Consequently, in analyzing messages in natural language and extracting the full meaning of a sentence, our brain separates sentences into words, identifies these, looks for specific key words, notes their order in the sentence, determines their meaning in context, and, most importantly, takes into account the content of the conversation.

A Language is not just a list of words, each language has its internal structure, generally called grammar. Whereas adhering to the grammar of a computer language is mandatory, natural language is much more flexible. In the understanding of natural language we may consider two levels of grammar. First there is the semantic grammar that deals with modifying of words to convey additional information e.g. "show", "Shows", "showing", "showed" and "shown", or "dog", "dogs", and "dog's". Second is the syntactic grammar which is associated with the structure of sentences. The latter grammar distinguishes between nouns, verbs, adjectives, adverbs, etc., allowing us to determine if "show" in a given sentence is a verb or a noun. When we write prose we try to adhere to grammatical rules and make the written information as explicit and unambiguous as possible. Under those condition we do not communicate through really unrestricted natural language since our language is restricted by grammatical rules. However, in our verbal communication we often break many of these rules with minimal loss of information; we can understand even poorly structured sentences based on their contextual content. We do not know how we, humans, do all this in our mind.

The question is - if so, how can we teach this to a dumb machine?

How can we make a computer understand the meaning of a sentence? In "understanding" we denote the recognition of the meaning of the sentence with a high sensitivity (low probability to declare the answer meaningless) and high specificity (low probability to ascribe to the answer a different meaning than the respondent wanted it to have). It is relatively easy to accept any answer as a correct answer (100% sensitivity) at the cost of lowering the specificity to zero. Alternatively we can make the specificity 100% by specifying just one specific sentence as a correct answer. The sensitivity will be practically zero since there is very little chance that the respondent in UNL will give exactly the expected specified answer. The art of teaching the computer to understand UNL is to maximize both sensitivity and specificity, i.e., to optimize the recognition process.

The first step in any computerized UNL scheme is to parse the sentence into words. This is very easy to do from the programming standpoint. Once we have isolated the individual words in the sentence, there are principally two different strategies to teach the computer to "understand" the meaning of a sentence. One is to analyze the grammatical structure of the sentence and assign to each word its grammatical attribute before trying to extract the meaning. The second is to try to identify the words or phrases in context with the previous steps of the conversation, and if necessary, look for other words or phrases that qualify the former words and provide a more specific understanding of the answer.

Although it may look difficult, it is possible to program a computer to analyze the grammatical structure of any sentence. This can be done using any programming language but is probably most readily achieved with PROLOG. Since CASIP is essentially a high computer language designed to analyze verbal input, and since we were, obviously, familiar with it, we have done this using a modified version CASIP. We used as landmarks the unambiguous identity of prepositions, articles, and pronouns as well as of adjectives ending with "al" and adverbs ending with "lly". We also used a set of common adjectives and adverbs that do not follow these rules, to identify those grammatical forms and us them as landmarks. We also used the suffixes "ed" and "ing" (excluding common nouns with this suffix) to confirm the identity of verbs and of suffixes such as "ies" "tion", and "sion" to confirm nouns.

After several trials and corrections, we had a program that could correctly identify the grammatical components of over 90% of grammatically correct sentences. This still did not mean that we were much closer to understanding the information that these sentences conveyed. To understand the meaning we must recognized the individual verbs and nouns, and this cannot be done without knowledge of the context. Moreover, when this program was tested by medical residents in a clinical simulation environment, most of the answers given had incorrect grammatical structures, leading to imprecise assignments of grammatical attributes to the different components of the sentences. This experience demonstrated the main shortcoming of this strategy of understanding UNL - unlike when writing an essay, medical students and residents in a CAI environment often break grammatical rules when they type in answers on a keyboard just as they do when they talk. This does not typify only medical students, of course, but is true of human verbal communication in general. We all break rules of grammar in our conversations.

The other strategy to teach computers to "understand" the meaning of sentences, involves identification of words in context. Once the computer identifies words in context with the conversation, it can branch out to identify additional words in context with the former words. For instance, if the question was "What treatment would you order in this case?" and the user answered "In this case I would Like to administer lidocaine iv", the compute will first identify "administer" and eliminate a group of other potential treatments, then it will branch out and identify "lidocaine" as one of a group of drugs that might be administered, and finally, if necessary (not in this particular case) it will scan a short list of different routes of giving drugs and identify "iv". If it did not identify a plausible drug in the second stage it would ask "What do you wish to administer?" At the first stage it will also check if the sentence does not contain "no", "cannot" or "never" to assure that the respondent is not one of those who cannot give a straight answer. The program would handle the answer in the same manner if instead of "administer" words such as "infuse", "inject", "given", or even "push" were used. All these words are contextual synonyms. This multistage strategy is more sophisticated than the simple hunt for "key words". (12)

The computer would handle the situation in the same manner even if the answer contained only two words, for example, "Infuse lidocaine".

"Lidocaine infusion" would also be accepted, as would be "Lidocaine iv". We may instruct the computer to accept also one-word sentences, e.g., "Lidocaine", but we prefer it to ask the student to repeat the answer in a more comprehensive manner. Such one or two word answers can be conversationally legitimate but they are grammatically awkward and could not be analyzed unambiguously by a grammatical structure analyzer. It is noteworthy that key words for syntactic analysis, such as "the", "in", "this", "I", "would" or "to" are ignored in the contextual analysis. CASIP skips these words outright and does not analyze them, thus shortening the time of sentence analysis.

The sequential recognition of words which takes place by repeated matching of a given word with many groups of contextual synonyms in the computer's memory is a time consuming process. This process an be started as soon as the respondent finishes the input, or it can be expected as soon as each word is typed in. We tried both approaches and found the latter to be operationally less effective and more time consuming, because the necessary contextual information is missing when the words are analyzed in a "half baked" answer.

The effectiveness of CASP's UNL recognition is somewhat limited by spelling mistakes of the respondents. This problem can be diminished but not eliminated by using for analysis words that are truncated (comparing, for example, just the first 5 letters of certain words in the answer with the words in the computer memory), or by placing a wild card letter in a position that is likely to be misspelled. This could be done to an extent that would not affect the specificity of word recognition. A more sophisticated approach was to include in the list of synonyms also phonetic synonyms to catch some awkward misspellings.

FROM TEACHING THE COMPUTER TO UNDERSTAND HUMAN LANGUAGE TO COMMUNICATION BETWEEN HUMANS

All this programming experience gave us material for thought about the way we humans communicate. Just like CASIP we understand each other in conversational UNL disregarding most grammatical constraints. Just like CASIP we store our information in groups of logical contextual synonym.(13) And, just like CASIP we anticipate contextual answers to our questions. If an answer is not contextual, we usually repeat our question. Occasionally we may decide that our conversant is not coherent. Now we must come up with answers to three questions: 1. Does this analogy between man-machine and man-man communication exist because we have programmed CASIP to follow our own pattern of communication? 2. Since it is assumed that mental processes proceed in parallel, while CASIP's understanding of UNL is a linear process, doesn't this prove that this analogy is coincidental? 3. Dies this computer based model represent a referred way to the understanding of natural language by both humans and machines?

Let us discuss the first point. When we developed CASIP's UNL capabilities we did not consciously follow any mental model. The objective was to achieve maximal "understanding" by the computer in minimum time of computation. Computers perform many human-like functions quite differently from humans. (14) There was no a priori reason, therefore, that our UNL recognition program should follow a model of human behavior. We tried different methods of sentence analysis and found the branched contextual synonym identification approach to be the most effective in terms of precision and speed. This does not mean that CASIP has reached the ultimate level of effectiveness of UNL recognition, ut following a Darwinian "natural selection" process it has emerged as best fit for the UNL CAI environment. Only after this mode of operation was selected, we realized that it resemble human behavior.

The second question relates to the assumption that parallel processing prevails in the human mind. This may be true though it is difficult to prove. (14) Parallel processing models change little conceptually in the way humans or computers behave. Even if it played a role in human verbal interaction, parallel processing is not unique to the human mind. In fact, it has been recently implemented in experimental computers. (15) Let us then consider the effect of parallel processing on the acceptance of UNL by computers. Based on our experience with CASIP, if real parallel processing (not interrupted serial multitasking) was routinely available on microcomputers we would attempt to recognize each word as soon as it is typed and keep accumulating all the words in the sentence for contextual analysis in parallel. If the anticipated answer had only a single key word or key phrase (comprising of two successive words) and we were ready to forego the risk of unexpected negation in the analyzed sentence, the computer could respond as soon as that key word was identified. This could increase the speed of recognition by a factor of a hundred. If more than one key word was required to convey the specific information we

expect, we would wait until the end of the message and then in parallel recognize each word in the answer, cutting down the time of this slow step in the UNL recognition process by an order of magnitude. In the latter case we would still have to process the whole sentence serially, since we must take into account the intrinsic contextual relationships. Considering mental parallel processing, it is conceivable that a routine analogous to this hypothetical one, which involves contextual key words, takes place in our minds. Just like our original serial model, this model involves the recognition of key words and the analysis of their contextual meaning. It does not negate, therefore, our basic assumption that one may use a serial computer program as a crude model for a mental process.

Finally, to answer the third question - we believe that the process we empirically developed for CASIP is fairly close to our own mental process. This assumption is not arbitrary. It is based on observations of human communication. Let us remember that a brief often non-grammatical answer evokes a faster response in humans just as it does in computers. Also remember that young children begin to communicate by single word sentences. Just like our medical residents, described above, children express a whole sentence through a single word. When a toddler says "Daddy" he may mean "Here is Daddy" or "I want Daddy to hold me". "Daddy, you are hurting me by pulling my arm", or "Daddy, I am hungry". The message, which is modulated by the intonation and supplemented by non verbal expressions, is usually understood in context with the situation. Messages with less associations than "Daddy" such as "Book" or "Spoon", which may mean "Here is a book" or "Give me this book" require minimal supplemental information except for the context of the situation, very much like "Lidocaine" cited above. As children develop they advance to two-word and then to three-word sentences. They manage to communicate effectively with their parents and with kids their age because these people anticipate a limited number of key words.

When we communicate with young children they, again, grope for key words; every nursery school teacher knows that to maintain children's attention you must use short sentences rich in key words. When young children read, they often misread non-key words, they add, skip or exchange prepositions, pronouns, adjectives or adverbs, but not the nouns - the subjects of the story, or the verbs, if it is important to know what the subjects are doing. If you come to think of it, we adults do the same when we skim over a newspaper; we look for key words just as a computer would do given the same task. If the computer is programmed to look for just a single word or a phrase, it will respond as soon as that word has been recognized, without waiting for the rest of the input. A short tempered person will do the same in conversation with you, responding even before you had a chance to finish speaking. A polite person will not cut into your words but may use the spare time to relax his mind. In either case the recipient of the message stops listening to you once the anticipated key words have been received. However, unlike computers, people may get bored while waiting for the next message, and their mind may start wandering, so they may miss the next message when it eventually comes. This is the reason that good orators, like good nursery school teachers, try to include just one piece of information in each sentence. Analogously, CASIP performs best when each answer contains just one message.

If human understanding of verbal communication is based on the recognition of key words just as an "intelligent" computer does, can we use the computer's UNL interactive program as a model for the study of impairments in human communication? Let us review the model to clarify the question. Following our model, verbal communication between people has the underlying expectation of a answer or statement that are within the contextual framework of the conversation. If the message is outside those limits, i.e. outside the range of probable answers or statements, we are unprepared and caught by surprise. If we cannot rationally place the message within the contextual framework, even after requesting further clarification, we decide that we must be conversing with an incoherent person. In expecting a message within a range of alternatives, each alternative constitutes of a group of logical contextual synonyms. For instance, even if we ask a general question such as "What do you want to do this Sunday morning?" we expect a large but still limited number of alternative responses. These may include, for example, go to church, shopping, jogging, reading, watching TV, mowing the lawn, Fixing the car, going to the beach, etc. Each of these alternative contextual messages comprises of a roup of messages which the listener considers interchangeable at that point in time. For example, "shopping", "purchasing","buying", "mall", "look for", "Sears", "K-Mark", "sale", "finding", are contextually equivalents in the expected answers. Any of these key words once recognized may be followed by "Don't you have anything better to do", "Didn't you say you wanted to fix the car this weekend?", or "Let us

go to the beach instead".

This model predicts impairments in two functional areas. One is a limited mental ability to come up with a sufficiently large assortment of contextual expectations. The other is insufficient or incoherent clustering of keywords for each of these expectations. The other is insufficient or incoherent clustering of keywords for each of these expectations. The number of contextually accepted alternative messages may be a measure of the brightness of the listener may be a measure of the brightness of the listener. A bright person is less likely to be surprised by an esoteric but still logical message. If the number of alternative anticipated answers is exceptionally small, we will say that the listener is dumb or single-track mined. However, if the expectations are numerous but consistently do not match the message because they all are out of field, the listener may have a serious cognitive problem.

Considering defects in clustering, we may encounter two types: Scanty clustering and incoherent clustering. Scanty clustering means an abnormally small number of contextual synonyms in each cluster. This may be the result of linguistic poverty, which can often be corrected, or it may arise from an inability to generalize or cluster contextually equivalent words, An inability to cluster contextual synonyms might result in an overload of expected alternative contextual messages. This would slow down the understanding of messages, impairing the flow of conversation, or the listener might even stop temporarily accepting messages, breaking the conversation altogether. This latter mechanisms may explain certain types of concentration deficiency, especially in children. Remedial action may be to give such children exercises in association or nouns, verbs and concepts.

Incoherent clustering may be a different cognitive defect. If certain clusters contain key words that do not belong there, one might encounter strange miscomprehensions. The listener might then respond in a way that shows that the message was misinterpreted. After confirming that such misunderstandings are not the result of semantic difficulties, which may be common in certain sub-populations such as normal bilingual children, one might elicit from the subject the rationale for the unexpected answer, which may be the result of an unconventional yet sound reasoning. If that attempt fails we may be confronted with an intrinsic mental problem.

If our model has merit in cognitive psychology how can it be tested? Moreover, how can the hypothetical defects cited above be diagnosed? We believe that both questions can be answered by measuring the time of response by either EEG (evoked multifocal response) or by EMG (picking up the action potential associated with speech even before utterance). The time of message processing may be defined as the time from the transmission of a key word to the time of response. If we use a computer to transmit the verbal message either visually on the screen or in an auditory mode, using a speech synthesizer, the transmission of key words can be clocked precisely. It should be possible to measure and characterize the delay in the response when an unexpected message is received, or the lack of response if attention (expectation of message) has ceased. Also the recovery time, the time it takes to be prepared for the next message (resetting the clusters) may be measurable by intentionally overloading the input, firing one message after the other as is done in the notorious cross examinations.

Children are often identified as having deficient communications skills. It is important to diagnose early the nature of their deficiency, and take remedial action as soon as practical. If the model we propose has merit, even if it is just a very crude approximation of the mental processes it tries to describe, it may lead to a more precise and quantitative assessment of some juvenile communication deficiencies.

In summary, we have discussed different ways to make a computer respond to unrestricted natural language and identified an effective mode of handling UNL by a computer. We then suggested this mode as a model for human verbal communication, Identified its elements and possible areas of impaired function. Finally we suggested an approach to diagnose these defects. The use of our computer-based experience in explaining certain aspects of human communication must be a gross oversimplification, but if this model will stimulate psychologists to consider such a possibility, test it and perhaps use it as a diagnostic tool, we may have made a modest contribution to that field.

REFERENCES

(1) Anbar M., "The Socratic Teacher - Computer Assisted Self Instruction," In M. Anbar (ed.) Computers in Medicine, Computer Science Press, Rockville, 1987 p. 203.

(2) Anbar M., and Loonsk, J.W., "Computer emulated oral exams - rationale and implementation of cue-free interactive computerized tests." MEdical Education (submitted), 1988.

(3) Anbar M., "Authoring languages for computer assisted instruction in medicine" CRC Critical Reviews in Medical Informatics, 1988.

(4) Anbar M., "Open-ended Computer Based Medical Instruction - Avoiding the Pitfalls of Multiple Choice Behavior," Resident & Staff Physician, vol.32, no.7, p.91, 1986.

(5) Abdulla A.M., Watkins L.O. and Henke J.H., "The Use of Natural Language Entry and Laser Disk Technology in CAI," J Med Educ vol. 59, p.739, 1984.

(6) Abdulla A.M., Henke J.S. and Walkins L.O., "Computer-Aided Learning: Experiences, Perspectives and Promises," J Am Coll cardiol vol.9, p.678, 1987.

(7) Anbar M., "CASIP - A Novel Authoring Tool for Open-ended Natural Language CAI," Proc. ADCIS (Association for Development of Computer-based Instruction Systems) 28th International Conf., ADCIS 1986, pp 1-4.

(8) Anbar M. "Computer Assisted Instruction (CAI), "A Way To Avoid The Pitfalls of Multiple Choice Behavior in Medical Practice," Medical Electronics vol. 18 no.4 p.118, 1987.

(9) Sylvester J.M. and Gardner D.W., "A Natural Language Interpreter for Reading Insulin Orders," Proc. AAMSI Sixth Annual Joint National Congress, Levy A.H. and Williams B.T. (Eds), 1987, p.283.

(10) Anbar M., "Socratic CAI _ Philosophy and Experience, "Proc. AAMSI Sixth Annual Joint Natl Congress, Levy A.H. & Williams B.T. (Eds), 1987, p. 286.

(11) Gardner A., Davidson J. and Barr A., "Understanding of Natural Language," Ch 4, in Handbook of Artificial Intelligence, A. Barr and E.A. Feigenbaum Eds, W. Kaufmann Inc., San Francisco 1981.

(12) Weizenbaum J., "ELIZA - A Computer Program for the Study of Natural Language Communications between Man and Machine," Comm. ACM vol. 9, no.1, p.36, 1966.

(13) Vellutino F.R., "dyslexia," Sci. Amer. vol.256 no.3, 34, 1987.

(14) Anbar M. and F. M. Snell, "Artificial Intelligence and it Impact on Medicine - Can the Computer Replace the Doctor," In M. Anbar (ed.) Computers in MEdicine, Computer Science Press, Rockville, 1987, p.280.

(15) Gelernter N., "Programming for Advanced Computing," Sci. Amer. vol.257, no.4, p.90, 1987.